

IN THE CLAIMS:

- 1 1. (Currently Amended) A system adapted to distribute redundant information across
2 disks of an array, the system comprising:
3 a storage operating system configured to invoke storage operations executed by a
4 storage system, the storage operating system further configured to manage storage of in-
5 formation, including the redundant information and data, on blocks of the disks in re-
6 sponse to disk access operations, the storage operating system including a storage module
7 adapted to compute the redundant information in response to placement of the data in
8 stripes across the disks, the storage operating system maintaining at least one unallocated
9 block per stripe for use by the storage module to store the computed redundant informa-
10 tion, wherein the at least one unallocated block used to store the redundant information is
11 located in any disk and wherein the location of the at least one unallocated block used to
12 store the redundant information is dynamically allocated by the storage module before
13 each write request for each stripe in a non-fixed pattern.
- 1 2. (Original) The system of Claim 1 wherein the storage module is a disk array controller
2 configured to compute the redundant information and reconstruct blocks lost due to fail-
3 ure of one or more of the disks.
- 1 3. (Original) The system of Claim 1 wherein the storage module is a RAID system con-
2 figured to compute the redundant information and reconstruct blocks lost due to failure of
3 one or more of the disks.
- 1 4. (Original) The system of Claim 3 wherein the storage operating system is further con-
2 figured to implement a high-level module that maintains information about locations of
3 the data on the disks.

- 1 5. (Original) The system of Claim 4 wherein the high-level module is a file system or da-
2 tabase adapted to control layout of the data on the disks.
- 1 6. (Original)The system of Claim 5 wherein the storage operating system integrates the
2 file system or database with the RAID system.
- 1 7. (Original)The system of Claim 6 wherein the file system or database is configured to
2 determine block locations of the data on the disks and the RAID system is configured to
3 determine the block locations of the redundant information on the disks.
- 1 8. (Original)The system of Claim 6 wherein the file system or database is configured to
2 determine block locations of the data and the redundant information on the disks.
- 1 9. (Original) The system of Claim 8 wherein the file system or database renders balanc-
2 ing decisions to determine the block locations of the data and the redundant information
3 on the disks.
- 1 10. (Original) The system of Claim 9 wherein the balancing decisions comprises one of
2 different sizes of disks, different speeds of disks, and whether a disk is more heavily util-
3 ized than other disks.
- 1 11. (Original) The system of Claim 8 further comprising block allocation map structures
2 used by the file system to determine the block locations of the data and the redundant in-
3 formation on the disks.
- 1 12. (Original)The system of Claim 11 wherein the redundant information is parity.

1 13. (Cancelled)

1 14. (Original) The system of Claim 1 wherein the storage module selects the at least one
2 unallocated block to store the redundant information and wherein the storage module
3 computes the redundant information using a redundant storage algorithm.

1 15. (Original) The system of Claim 14 wherein the selection of the at least one unallo-
2 cated block to store redundant information is independent of the redundant storage algo-
3 rithm.

1 16. (Previously Presented) The system of Claim 15 wherein the redundant storage algo-
2 rithm is a symmetric algorithm.

1 17. (Original) The system of Claim 16 wherein the redundant information is parity.

1 18. (Original) The system of Claim 14 wherein the at least one unallocated block used to
2 store the redundant information comprises two or more unallocated blocks used to store
3 the redundant information.

1 19. (Original) The system of Claim 18 wherein the selection of the unallocated blocks to
2 store redundant information is independent of the redundant storage algorithm used to
3 compute the redundant information.

1 20. (Original) The system of Claim 19 wherein the redundant storage algorithm depends
2 on positions of the blocks in the array.

1 21. (Previously Presented) The system of Claim 20 wherein the redundant storage algo-
2 rithm is a symmetric and wherein the redundant information is parity.

1 22. (Currently Amended) A method for distributing redundant information across disks of
2 an array, the method comprising the steps of:

3 dividing each disk into blocks, the blocks being organized into stripes such that
4 each stripe contains one block from each disk;

5 selecting any blocks in the stripe not used to contain data to contain the redundant
6 information, wherein the block used to contain the redundant information is located in
7 any disk and wherein the location of the block used to contain the redundant information
8 is dynamically allocated before each write request for each stripe in a non-fixed pattern;
9 and

10 computing the redundant information based on contents of all other blocks in the
11 stripe, regardless of whether the blocks contain data.

1 23. (Original)The method of Claim 22 further comprising the step of determining which
2 block in a stripe contains redundant information each time there is a write request to the
3 stripe.

1 24. (Original)The method of Claim 23 further comprising the step of assigning a block to
2 contain redundant information when each stripe is written.

1 25. (Original) The method of Claim 24 wherein the step of determining is performed by a
2 high-level module of a storage system and wherein the steps of computing and assigning
3 are performed by a storage module of the storage system.

1 26. (Original) The method of Claim 25 further comprising the steps of:

2 maintaining, by the high-level module, at least one unallocated block per stripe
3 for use by the storage module; and
4 providing an indication from the high-level module to the storage module of the
5 unallocated block to contain parity.

1 27. (Original) The method of Claim 26 further comprising the step of reconstructing, us-
2 ing the storage module, a block that is lost due to failure of a disk.

1 28. (Original) The method of Claim 25 wherein the high-level module is a file system and
2 wherein the storage module is one of an array controller and a RAID system.

1 29. (Original) The method of Claim 28 wherein the step of computing comprises the step
2 of computing the redundant information in response to placement of the data in stripes
3 across the disks.

1 30. (Original) The method of Claim 29 wherein the step of computing further comprises
2 the step of computing the redundant information using algebraic and algorithmic calcula-
3 tions in response to the placement of the data on the array.

1 31. (Currently Amended) Apparatus for distributing redundant information across disks
2 of an array, the apparatus comprising:

3 means for dividing each disk into stripes, with each stripe containing one block
4 from each disk;

5 means for selecting any blocks in the stripe not used to contain data to contain re-
6 dundant information, wherein the block used to contain the redundant information is lo-
7 cated in any disk and wherein the location of the block used to contain the redundant in-
8 formation is dynamically allocated before each write request for each stripe in a non-
9 fixed pattern; and

10 means for computing the redundant information based on contents of all other
11 blocks in the stripe, regardless of whether the blocks contain data.

1 32. (Original) The apparatus of Claim 31 further comprising means for determining
2 which block or blocks in a stripe holds redundant information each time there is a write
3 operation to the stripe.

1 33. (Currently Amended) A computer readable medium containing executable program
2 instructions for distributing parity across disks of an array, the executable instructions
3 comprising one or more program instructions for:
4 dividing each disk into stripes, with each stripe containing one block from each
5 disk;
6 selecting any blocks in the stripe not used to contain data to contain parity,
7 wherein the block used to contain the parity is located in any disk and wherein the loca-
8 tion of the block used to contain the parity is dynamically allocated before each write re-
9 quest for each stripe in a non-fixed pattern; and
10 computing the parity based on contents of all other blocks in the stripe, regardless
11 of whether the blocks contain data.

1 34. (Previously Presented) The system of Claim 15 wherein the redundant storage algo-
2 rithm is an asymmetric algorithm

1 35. (Previously Presented) The system of Claim 20 wherein the redundant storage algo-
2 rithm is an asymmetric algorithm and wherein the redundant information is parity.

1 36. (Previously Presented) A method for distributing redundant information across disks
2 of an array with a plurality of blocks on each disk, comprising:
3 determining which blocks are unallocated in a stripe across the disks;
4 reserving unallocated blocks for storing the redundant information in one or more
5 reserved unallocated blocks;
6 arranging data in the stripe for the data to be stored in one or more allocated
7 blocks across the disks of the array;
8 assigning the redundant information to the one or more reserved unallocated
9 blocks; and
10 writing the data in the allocated blocks and the redundant information in the one
11 or more reserved unallocated blocks as the stripe across the disks of the array.

1 37. (Previously Presented) The method of claim 36, further comprising:
2 storing parity information as the redundant information in the one or more re-
3 served unallocated blocks.

1 38. (Previously Presented) The method of claim 36, further comprising:
2 adding a disk to the array;
3 storing a second stripe across the array by determining one or more unallocated
4 blocks across the array including the added disk, and writing the data to allocated blocks
5 and the redundant information to the one or more unallocated blocks of the second stripe.

1 39. (Previously Presented) An apparatus for distributing redundant information across
2 disks of an array with a plurality of blocks on each disk, comprising:
3 means for determining which blocks are unallocated in a stripe across the disks;
4 means for reserving unallocated blocks for storing the redundant information in
5 one or more reserved unallocated blocks;
6 means for arranging data in the stripe for the data to be stored in one or more allo-
7 cated blocks across the disks of the array;
8 means for assigning the redundant information to the one or more reserved unal-
9 located block; and
10 means for writing the data in the allocated blocks and the redundant information
11 in the one or more reserved unallocated blocks as the stripe across the disks of the array.

1 40. (Previously Presented) The apparatus of claim 39, further comprising:
2 means for storing parity information as the redundant information in the one or
3 more reserved unallocated blocks.

1 41. (Previously Presented) The apparatus of claim 39, further comprising:
2 means for adding a disk to the array;
3 means for storing a second stripe across the array by determining one or more un-
4 allocated blocks across the array including the added disk, and writing the data to allo-
5 cated blocks and the redundant information to the one or more unallocated blocks of the
6 second stripe.

1 42. (Previously Presented) A system for distributing redundant information across disks
2 of an array with a plurality of blocks on each disk, comprising:
3 a storage operating system configured to invoke storage operations executed by a
4 storage system, the storage operating system further configured (i) to determine which
5 blocks are unallocated in a stripe across the disks, (ii) to reserve unallocated blocks for
6 storing the redundant information in one or more reserved unallocated blocks, (iii) to ar-
7 range data in the stripe for the data to be stored in one or more allocated blocks across the
8 disks of the array, (iv) to assign the redundant information to the one or more reserved
9 unallocated block, and (iv) to write the data in the allocated blocks and the redundant in-
10 formation in the one or more reserved unallocated blocks as the stripe across the disks of
11 the array.

1 43. (Previously Presented) The system of claim 42, wherein the redundant information is
2 parity information.

1 44. (Previously Presented) A computer readable medium containing executable program
2 instructions for distributing parity across disks of an array, the executable instructions
3 comprising one or more program instructions for:
4 determining which blocks are unallocated in a stripe across the disks, where each
5 disk has a plurality of blocks;

6 reserving unallocated blocks for storing the redundant information in one or more
7 reserved unallocated blocks;
8 arranging data in the stripe for the data to be stored in one or more allocated
9 blocks across the disks of the array; and
10 assigning the redundant information to the one or more reserved unallocated
11 block; and
12 writing the data in the allocated blocks and the redundant information in the one
13 or more reserved unallocated blocks as the stripe across the disks of the array.